

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Appl. No. : 10/815,233 Confirmation No. 8121
Applicants : Chen et al.
Title: Method Enabling Real-Time Testing of On-Demand Infrastructure to Predict Service Level Agreement Compliance
Filed : March 31, 2004
TC/A.U. : 2123
Examiner : Mary C. Jacob
Docket No. : AUS920031048US1
Customer No. : 48,916

BRIEF ON APPEAL

Mail Stop Appeal Briefs-Patents
Commissioner for Patents
P.O. Box 1450, Alexandria, VA 22313-1450

This is an appeal from the Final Office Action, mailed on September 22, 2009, finally rejecting claims 1-8, 10-18, 20-25 and 27.

The fee for filing this Brief on Appeal is \$540.00 and is being paid electronically at the time of filing of this Brief. A Request for a Three-Month Extension of Time to file this Amendment by July 22, 2009 is also being filed and paid for concurrently with this filing. If these amounts are insufficient, or should any additional fees under 37 C.F.R. § 1.16 to 1.21 be required for any reason relating to the enclosed materials, the Commissioner is authorized to deduct said fees from IBM Corporation, Deposit Account No. 09-0457.

TABLE OF CONTENTS

REAL PARTY IN INTEREST	3
RELATED APPEALS AND INTERFERENCES	4
STATUS OF CLAIMS	5
STATUS OF AMENDMENTS	6
SUMMARY OF CLAIMED SUBJECT MATTER	7-13
GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL	14
ARGUMENTS OF APPELLANTS	15-21
CLAIMS APPENDIX	22-29
EVIDENCE APPENDIX	30
RELATED PROCEEDINGS APPENDIX	31

REAL PARTY IN INTEREST

The real party in interest is International Business Machines Corp. (IBM) of Armonk, New York

RELATED APPEALS AND INTERFERENCES

The Appellants are not aware of any related appeals, interferences or judicial proceedings that will directly affect, be directly affected by or have a bearing on the Board's decision in the pending appeal.

STATUS OF CLAIMS

Claims 1-8, 10-18, 20-25 and 27 are currently pending and have been finally rejected.

Claim 14 is objected to because of informalities. The objection of claim 14 is not being appealed but will be addressed in the event the §103 rejection of claim 14 is reversed.

Claims 1-4, 7, 8, 10, 12-18, 20-25 and 27 are rejected under 35 U.S.C. §103(a) as being unpatentable over Chandra et al. ("An Online Optimization-based Technique for Dynamic Resource Allocation in GPS Servers," Technical Report UM-CS-2002-030, University of Massachusetts, July 2002; hereinafter referred to as "Chandra") in view of D'Arienzo et al. ("Automatic SLA Management in SLA-Aware Architecture," 10th International Conference on Telecommunications, 23 Feb. – 1 Mar. 2003, Vol. 2, pp. 1402-1406; hereinafter referred to as "D'Arienzo").

Claim 5 is rejected under 35 U.S.C. §103(a) as being unpatentable over Chandra in view of D'Arienzo and further in view of Chan (U.S. Pat. No. 6,466,898).

Claim 6 is rejected under 35 U.S.C. §103(a) as being unpatentable over Chandra in view of D'Arienzo and further in view of Sheets et al. (U.S. Pat. No. 6,816,905; hereinafter referred to as "Sheets").

Claim 11 is rejected under 35 U.S.C. §103(a) as being unpatentable over Chandra in view of D'Arienzo and further in view of Nargarajan et al. ("Modelling and Simulation on Alarm Base Network Management System for Effective SLA Monitoring and Management," SCI 2003, 7th World Multiconference on Systemics, Cybernetics and Informatics Proceedings, Jul. 27-30, 2003; hereinafter referred to as "Nargarajan").

The rejections of claims 1-8, 10-18, 20-25 and 27 are being appealed.

STATUS OF AMENDMENTS

All amendments to the claims have been entered.

SUMMARY OF CLAIMED SUBJECT MATTER

In general, the present invention relates to a system and method for predicting resource allocation in an on-demand services environment and generating a service level agreement (SLA) consistent with the prediction. Each claim being appealed is summarized below. References to the Specification refer to the paragraphs of the published application, U.S. Pat. App. No. 2005/0222885, published October 6, 2005.

Claim 1 is an independent method claim for predicting a service level in a utility computing environment (100, FIG. 1; ¶¶[0020]-[0022]). The method includes creating a resource profile (201 & 211, FIG. 5; ¶0041) corresponding to a dynamically allocated subset of computing resources (151-158; FIG. 2; ¶[0026]) allocated according to a service level agreement (153, 161; FIG. 2; ¶[0028]); loading (225, 161; FIG. 6; ¶[0052]) a workload profile (213; FIG. 5; ¶[0042]) representing a demand profile (¶[0042]) for an enterprise (101; FIG. 1; ¶[0020]); simulating processing of the workload profile (229; FIG. 6; ¶[0054]) using the resource profile (201 & 211, FIG. 5; ¶0041) to produce a service level result (¶[0049]), wherein the resource profile (201 & 211, FIG. 5; ¶0041) subset is modified (227; FIG. 6; ¶¶[0053]-[0054]) during the simulation according to the service level agreement (235; FIG. 6; ¶[0055]); and generating a new service level agreement (153, 161; FIG. 2; ¶[0028]) in the event the resource profile (153, 161; FIG. 2; ¶[0028]) cannot process the workload profile (229; FIG. 6; ¶[0054]) at an expected service level corresponding to the service level agreement (153, 161; FIG. 2; ¶[0028]), wherein the new service level agreement (153, 161; FIG. 2; ¶[0028]) will process the workload profile (229; FIG. 6; ¶[0054]) at the expected service level.

Claim 2 is a dependent claim that includes all the elements of **Claim 1**, as described above, for predicting a service level in a utility computing environment (100, FIG. 1; ¶¶[0020]-[0022]). In addition, **Claim 2** includes comparing (231; FIG. 6; ¶[0054]) the service level result (¶[0049]) to the service level agreement (153, 161; FIG. 2; ¶[0028]); and signaling (¶[0050]) whether the resource profile (153, 161; FIG. 2; ¶[0028]) will process the workload profile (229; FIG. 6; ¶[0054]) at an expected service level corresponding to the service level agreement (153, 161; FIG. 2; ¶[0028]).

Claim 3 is a dependent claim that includes all the elements of **Claim 1**, as described above, for predicting a service level in a utility computing environment (100, FIG. 1; ¶¶[0020]-[0022]). In addition, **Claim 3** includes the limitation that the subset of computing resources (151-158; FIG. 2; ¶[0026]) includes allocated processing resources (¶[0022]) and memory resources (¶[0022]) for a client account.

Claim 4 is a dependent claim that includes all the elements of **Claim 1**, as described above, for predicting a service level in a utility computing environment (100, FIG. 1; ¶¶[0020]-[0022]). In addition, **Claim 4** includes the limitation that the service level agreement (153, 161; FIG. 2; ¶[0028]) includes a base resource allocation (¶[0041]), a maximum resource allocation (¶[0041]), resource costs (¶[0041]), and rules for dynamically reallocating the resources based upon workload demand (¶[0041]).

Claim 5 is a dependent claim that includes all the elements of **Claim 1**, as described above, for predicting a service level in a utility computing environment (100, FIG. 1; ¶¶[0020]-[0022]). In addition, **Claim 5** includes the limitation that the simulation is scheduled to run automatically at an off-peak time (¶[0046]).

Claim 6 is a dependent claim that includes all the elements of **Claim 1**, as described above, for predicting a service level in a utility computing environment (100, FIG. 1; ¶¶[0020]-[0022]). In addition, **Claim 6** includes determining a cost associated with meeting the new service level agreement (¶[0055]).

Claim 7 is a dependent claim that includes all the elements of **Claim 1**, as described above, for predicting a service level in a utility computing environment (100, FIG. 1; ¶¶[0020]-[0022]). In addition, **Claim 7** includes the limitation that the resource profile (153, 161; FIG. 2; ¶[0028]) includes a communication bandwidth allocation (¶[0041]).

Claim 8 is a dependent claim that includes all the elements of **Claim 1**, as described above, for predicting a service level in a utility computing environment (100, FIG. 1; ¶¶[0020]-

[0022]). In addition, **Claim 8** includes comparing the workload profile (213; FIG. 5; ¶[0042]) to a second workload profile representing an actual demand profile for a second client account; wherein the simulating is based upon a result of the comparison (¶[0047]).

Claim 10 is a dependent claim that includes all the elements of **Claim 1**, as described above, for predicting a service level in a utility computing environment (100, FIG. 1; ¶¶[0020]-[0022]). In addition, **Claim 10** includes the limitation that the workload profile (213; FIG. 5; ¶[0042]) includes scheduling information (¶[0046]) and the simulation step incorporates the scheduling information in the processing.

Claim 11 is a dependent claim that includes all the elements of **Claim 1**, as described above, for predicting a service level in a utility computing environment (100, FIG. 1; ¶¶[0020]-[0022]). In addition, **Claim 11** includes the limitation that the workload profile (213; FIG. 5; ¶[0042]) includes information corresponding to one or both of prioritization of resources and importance of specific resources (¶[0052]).

Claim 12 is a dependent claim that includes all the elements of **Claim 1**, as described above, for predicting a service level in a utility computing environment (100, FIG. 1; ¶¶[0020]-[0022]). In addition, **Claim 12** includes the limitation that the workload profile is loaded from a configuration file (¶[0053]).

Claim 13 is an independent apparatus claim for a system that predicts a service level in a utility computing environment (100, FIG. 1; ¶¶[0020]-[0022]). The system includes an allocated subset of the set of computing resources (151-158; FIG. 2; ¶[0026]) and logic (¶[0019]) for loading a workload profile (213; FIG. 5; ¶[0042]) representing a hypothetical demand profile (¶[0042]) for a client account; for simulating processing of the workload profile (229; FIG. 6; ¶[0054]), wherein the workload profile (213; FIG. 5; ¶[0042]) is based upon actual, measured data (¶[0043]), using the allocated subset of the set of available computing resources (151-158; FIG. 2; ¶[0026]) to produce a service level result (¶[0049]); for modifying the allocated subset of the available computing resources (151-158; FIG. 2; ¶[0026]) based upon the service level result (¶[0049]); and for generating a new service level agreement (153, 161; FIG. 2; ¶[0028]) in

the event the simulation produced by the simulation logic cannot process the workload profile at an expected service level corresponding to the first service level agreement, wherein the new service level agreement (153, 161; FIG. 2; ¶[0028]) will process the workload profile at the expected service level (213; FIG. 5; ¶[0042]).

Claim 14 is a dependent apparatus claim that includes all the elements of **Claim 13**, as described above, for a system that predicts a service level in a utility computing environment (100, FIG. 1; ¶¶[0020]-[0022]). In addition, **Claim 14** includes logic for comparing (231; FIG. 6; ¶[0054]) the service level result (¶[0049]) to a second service level agreement (153, 161; FIG. 2; ¶[0028]); and for signaling (¶[0050]) whether or not modified, allocated subset of the available computing resources (151-158; FIG. 2; ¶[0026]) will process the workload profile (229; FIG. 6; ¶[0054]) at an expected service level corresponding to the second service level agreement (153, 161; FIG. 2; ¶[0028]).

Claim 15 is a dependent apparatus claim that includes all the elements of **Claim 13**, as described above, for a system that predicts a service level in a utility computing environment (100, FIG. 1; ¶¶[0020]-[0022]). In addition, **Claim 15** includes the limitation that the set of available computing resources (151-158; FIG. 2; ¶[0026]) include processing resources (¶[0022]); and memory resources (¶[0022]).

Claim 16 is a dependent apparatus claim that includes all the elements of **Claims 13 and 15**, as described above, for a system that predicts a service level in a utility computing environment (100, FIG. 1; ¶¶[0020]-[0022]). In addition, **Claim 16** includes the limitations that the set of available computing resources (151-158; FIG. 2; ¶[0026]) includes a base resource allocation (¶[0041]), a maximum resource allocation (¶[0041]), resource costs (¶[0041]); and rules for dynamically reallocating the resources based upon workload demand (¶[0041]).

Claim 17 is a dependent apparatus claim that includes all the elements of **Claims 13 and 15**, as described above, for a system that predicts a service level in a utility computing environment (100, FIG. 1; ¶¶[0020]-[0022]). In addition, **Claim 17** includes the limitation that the set of

available computing resources (151-158; FIG. 2; ¶[0026]) further comprises communication bandwidth (¶[0041]).

Claim 18 is a dependent apparatus claim that includes all the elements of **Claim 13**, as described above, for a system that predicts a service level in a utility computing environment (100, FIG. 1; ¶¶[0020]-[0022]). In addition, **Claim 18** includes logic for comparing the workload profile (213; FIG. 5; ¶[0042]) to a second workload profile representing an actual demand profile for a second client account, wherein a simulation produced by the simulation logic is based upon a result of the comparison step (¶[0047]).

Claim 20 is a dependent apparatus claim that includes all the elements of **Claim 13**, as described above, for a system that predicts a service level in a utility computing environment (100, FIG. 1; ¶¶[0020]-[0022]). In addition, **Claim 20** includes the limitation that the workload profile (213; FIG. 5; ¶[0042]) includes scheduling information (¶[0046]) and the simulation logic incorporates the scheduling information in the processing.

Claim 21 is an independent method of manufacture claim for a programming product that predicts a service level in a utility computing environment (100, FIG. 1; ¶¶[0020]-[0022]). The programming product includes a memory (161, FIG. 5; ¶[0019]), a resource list (201 & 211, FIG. 5; ¶[0041]) detailing a set of available computing resources (151-158; FIG. 2; ¶[0026]); an allocated resource list (201 & 211, FIG. 5; ¶[0041]) detailing an allocated subset of the set of available computing resources (151-158; FIG. 2; ¶[0026]); and logic (¶[0019]) for creating a computer resource profile (201 & 211, FIG. 5; ¶[0041]) based upon the allocated subset of the set of available computing resources; for loading a workload profile (213; FIG. 5; ¶[0042]) representing a hypothetical demand profile for a client account; for simulating the processing of the workload profile (229; FIG. 6; ¶[0054]) using the computer resource profile to produce a service level result (¶[0049]); for comparing the service level result (¶[0049]) to a service level agreement (153, 161; FIG. 2; ¶[0028]); for signaling (¶[0050]) whether the computing resource profile (201 & 211, FIG. 5; ¶[0041]) will process the workload profile (213; FIG. 5; ¶[0042]) at an expected service level corresponding to the service level agreement (153, 161; FIG. 2; ¶[0028]); and for generating a new service agreement (153, 161; FIG. 2; ¶[0028]) in the event

the computing resource profile (201 & 211, **FIG. 5**; ¶[0041]) cannot process the workload profile (213; **FIG. 5**; ¶[0042]) at the expected service level corresponding to the service level agreement (153, 161; **FIG. 2**; ¶[0028]), wherein the new service level agreement (153, 161; **FIG. 2**; ¶[0028]) will process the workload profile (229; **FIG. 6**; ¶[0054]) at the expected service level.

Claim 22 is a dependent method of manufacture claim that includes all the elements of **Claim 21**, as described above, for a programming product that predicts a service level in a utility computing environment (100, **FIG. 1**; ¶¶[0020]-[0022]). In addition, **Claim 22** includes the limitations that the computing resource profile includes processing resources (¶[0021]) and memory resources (¶[0022]).

Claim 23 is a dependent method of manufacture claim that includes all the elements of **Claims 21 and 22**, as described above, for a programming product that predicts a service level in a utility computing environment (100, **FIG. 1**; ¶¶[0020]-[0022]). In addition, **Claim 23** includes the limitations that the computing resource profile (153, 161; **FIG. 2**; ¶[0028]) includes a base resource allocation (¶[0041]), a maximum resource allocation (¶[0041]), resource costs (¶[0041]) and rules for dynamically reallocating the resources based upon workload demand (¶[0041]).

Claim 24 is a dependent method of manufacture claim that includes all the elements of **Claims 21 and 22**, as described above, for a programming product that predicts a service level in a utility computing environment (100, **FIG. 1**; ¶¶[0020]-[0022]). In addition, **Claim 24** includes the limitations that the computing resource profile (153, 161; **FIG. 2**; ¶[0028]) also comprises a communication bandwidth allocation (¶[0041]).

Claim 25 is a dependent method of manufacture claim that includes all the elements of **Claim 21**, as described above, for a programming product that predicts a service level in a utility computing environment (100, **FIG. 1**; ¶¶[0020]-[0022]). In addition, **Claim 25** includes logic (¶[0019]) for comparing the workload profile (213; **FIG. 5**; ¶[0042]) to a second workload profile (213; **FIG. 5**; ¶[0042]) representing an actual demand profile for a second client account,

wherein a simulation produced by the simulation logic is based upon a result of the comparison step (¶[0047]).

Claim 27 is a dependent method of manufacture claim that includes all the elements of **Claim 21**, as described above, for a programming product that predicts a service level in a utility computing environment (100, FIG. 1; ¶¶[0020]-[0022]). In addition, **Claim 27** includes the limitation that the workload profile (213; FIG. 5; ¶[0042]) includes scheduling information (¶[0046]) and the simulation logic incorporates the scheduling information in the processing.

GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

1. Whether a group of claims consisting of **independent claims 1, 13 and 21 and dependent claims 2-4, 7, 8, 10, 12, 14-18, 22-25 and 27** is unpatentable under 35 U.S.C. §103(a) over Chandra et al. ("An Online Optimization-based Technique for Dynamic Resource Allocation in GPS Servers," Technical Report UM-CS-2002-030, University of Massachusetts, July 2002; hereinafter referred to as "Chandra") in view of D'Arienzo et al. ("Automatic SLA Management in SLA-Aware Architecture," 10th International Conference on Telecommunications, 23 Feb. – 1 Mar. 2003, Vol. 2, pp. 1402-1406; hereinafter referred to as "D'Arienzo").
2. Whether a group of claims consisting of **dependent 5** is unpatentable under 35 U.S.C. §103(a) over Chandra in view of D'Arienzo and further in view of Chan (U.S. Pat. No. 6,466,898).
3. Whether a group of claims consisting of **dependent 6** is unpatentable under 35 U.S.C. §103(a) over Chandra in view of D'Arienzo and further in view of Sheets et al. (U.S. Pat. No. 6,816,905; hereinafter referred to as "Sheets").
4. Whether a group of claims consisting of **dependent 11** is unpatentable under 35 U.S.C. §103(a) over Chandra in view of D'Arienzo and further in view of Nargarajan et al. ("Modelling and Simulation on Alarm Base Nettwork Management System for Effective SLA Monitoring and Management," SCI 2003, 7th World Multiconference on Systemics, Cybernetics and Informatics Proceedings, Jul. 27-30, 2003; hereinafter referred to as "Nargarajan")

ARGUMENTS OF APPELLANTS

1. Whether a group of claims consisting of independent claims 1, 13 and 21 and dependent claims 2-4, 7, 8, 10, 12, 14-18, 22-25 and 27 is unpatentable under 35 U.S.C. §103(a) over Chandra et al. ("An Online Optimization-based Technique for Dynamic Resource Allocation in GPS Servers," Technical Report UM-CS-2002-030, University of Massachusetts, July 2002; hereinafter referred to as "Chandra") in view of D'Arienzo et al. ("Automatic SLA Management in SLA-Aware Architecture," 10th International Conference on Telecommunications, 23 Feb. – 1 Mar. 2003, Vol. 2, pp. 1402-1406; hereinafter referred to as "D'Arienzo").

Claims 1, 13 and 21

Basically, Chandra is directed to "online workload predictions and optimization-based techniques to dynamically allocate resources to competing web applications running on shared servers." (p. 1, Abstract, lines 8-11; *emphasis added*). In other words, Chandra addresses resource allocation in a shared computing environment. In contrast, Appellants' claimed subject matter is directed at generating a service level agreement rather than the actual allocation of resources while in the process of providing those resources. While Chandra addresses techniques for the reallocation of resources among competing users, perhaps for the purpose of providing resources to satisfy a particular agreement, there is no teaching or suggestion directed to modifying an agreement such as a SLA. D'Arienzo does not address this deficiency because, although D'Arienzo mentions that an agreement may be re-negotiated or replaced, neither D'Arienzo nor Chandra, either alone or in combination, describe any mechanism by which this can be accomplished.

In short, Chandra is directed to the monitoring and management of actual resources in a computing environment, not to the monitoring and management of a service level agreement (SLA). These are two unrelated technologies. Chandra does not even mention SLAs is therefore an inappropriate basis for a rejection of Appellants' claims. In addition, there is simply no expectation that Chandra would be combined with D'Arienzo by one with skill in the corresponding arts because, like Appellants' claimed subject matter, D'Arienzo is directed to an entirely different technology than Chandra. In other words, the current Office action (O.A.) is

simply combining two unrelated references, at least one of one of which is unrelated to Appellants' claimed technology, i.e. the monitoring and management of an SLA.

Specific elements of Chandra are mischaracterized to correspond to elements of Appellants' claimed subject matter. For example, one cited portion of Chandra states:

An alternative approach is to allocate resources to applications based upon the variations in their workloads. In this approach, each application is given a certain minimum share based upon coarse-grain estimates of its resource needs...

(p.1, c. 2, lines 18-29). This excerpt is mischaracterized to as a "resource profile" when, if anything, it seems to be directed more to the idea of a workload profile or a demand profile, although Appellants do not concede either. In fact, the O.A. seems to rely upon this excerpt for a "resource profile," a "workload profile" and a "demand profile." Clearly, it cannot simultaneously stand for all three. Appellants also contend neither D'Arienzo nor any of the other cited art provides these particular elements that Chandra lacks.

In response to these arguments, the Final Office Action, dated September 22, 2008, (FOA), states:

Applicant argues that Chandra is directed to the monitoring and management of actual resources in a computing environment, not to the monitoring and management of a service level agreement (page 11). The Examiner would like to note that the claims are directed to (claim 1) "a method of *predicting service level*", (claim 13) "a system for *simulating service* in a utility computing environment having a first service level agreement" and (claim 21), "a computer program product for *predicting service level* compliance in a utility computing environment having a service level agreement" ...

(p. 12, lines 3-12; citing Appellants' Amendment filed 07/13/08). Appellants submit that the reliance on the claim preambles is inappropriate. The M.P.E.P. states as follows:

The claim preamble must be read in the context of the entire claim. The determination of whether preamble recitations are structural limitations or mere statements of purpose or use "can be resolved only on review of the entirety of the [record] to gain an understanding of what the inventors actually invented and

intended to encompass by the claim." *Corning Glass Works*, 868 F.2d at 1257, 9 USPQ2d at 1966. If the body of a claim fully and intrinsically sets forth all of the limitations of the claimed invention, and the preamble merely states, for example, the purpose or intended use of the invention, rather than any distinct definition of any of the claimed invention's limitations, then the preamble is not considered a limitation and is of no significance to claim construction.

(M.P.E.P. §2111.02 II; citing *Pitney Bowes, Inc. v. Hewlett-Packard Co.*, 182 F.3d 1298, 1305, 51 USPQ2d 1161, 1165 (Fed. Cir. 1999); *emphasis added*). In other words, since the claim limitations clearly define that which the Inventors regarded as the invention, limitations should not be inferred from the claim preambles.

For the reasons stated above, Appellants respectfully request reversal of the §103(a) rejection of claims 1, 13 and 21.

Claims 8, 18 and 25

Claims 8, 18 and 25 are directed to "comparing the workload profile to a second workload profile... wherein the simulation is based upon the comparison." The cited portion of Chandra (p. 9, §C.2, ¶¶ 1 and 2) is comparing two different applications for the purpose of a potential resource allocation rather than to use as the basis of a simulation. Further, claims 8, 18 and 25 are each allowable because each depends upon one of the allowable independent claims.

For both the reasons stated above and those explained in conjunction with claims 1, 13 and 21, Appellants respectfully request reversal of the §103(a) rejection of claims 8, 18 and 25.

Claims 2-4, 7, 10, 12, 14-17, 22-24 and 27

Claims 2-4, 7, 8, 10, 12, 14-18, 22-25 and 27 are each allowable because each depends upon one of the allowable independent claims. For both the reasons stated above and those explained in conjunction with claims 1, 13 and 21, Appellants respectfully request reversal of the §103(a) rejection of claims 2-4, 7, 8, 10, 12, 14-18, 22-25 and 27.

2. Whether a group of claims consisting of dependent 5 is unpatentable under 35 U.S.C. §103(a) over Chandra in view of D'Arienzo and further in view of Chan (U.S. Pat. No. 6,466,898).

Claim 5 is allowable because it depends upon allowable independent claim 1. Further Chan does not suggest or teach any of the elements Appellants have argued are lacking in Chandra and D'Areinzo as explained above with respect to claims 1, 13 and 21. Therefore, for both the reasons stated above and those explained in conjunction with claims 1, 13 and 21, Appellants respectfully request reversal of the §103(a) rejection of claim 5.

3. **Whether a group of claims consisting of dependent 6 is unpatentable under 35 U.S.C. §103(a) over Chandra in view of D'Arienzo and further in view of Sheets et al. (U.S. Pat. No. 6,816,905; hereinafter referred to as "Sheets").**

Claim 6 is allowable because it depends upon allowable independent claim 1. Further, Sheets does not suggest or teach any of the elements Appellants have argued are lacking in Chandra and D'Areinzo as explained above with respect to claims 1, 13 and 21. Therefore, for both the reasons stated above and those explained in conjunction with claims 1, 13 and 21, Appellants respectfully request reversal of the §103(a) rejection of claim 6.

4. Whether a group of claims consisting of dependent 11 is unpatentable under 35 U.S.C. §103(a) over Chandra in view of D'Arienzo and further in view of Nargarajan et al. ("Modelling and Simulation on Alarm Base Nettwork Management System for Effective SLA Monitoring and Management," SCI 2003, 7th World Multiconference on Systemics, Cybernetics and Informatics Proceedings, Jul. 27-30, 2003; hereinafter referred to as "Nargarajan")

Claim 11 is allowable because it depends upon allowable independent claim 1. Further, Nargarajan does not suggest or teach any of the elements Appellants have argued are lacking in Chandra and D'Areinzo as explained above with respect to claims 1, 13 and 21. Therefore, for both the reasons stated above and those explained in conjunction with claims 1, 13 and 21, Appellants respectfully request reversal of the §103(a) rejection of claim 11.

CONCLUSION

To establish *prima facie* obviousness of a claimed invention under §103(a), all the claim limitations must be taught or suggested by the prior art. (M.P.E.P., §2143,03, citing *in re Royka*, 490 F.2d 981; 180 U.S.P.Q. 580 (CCPA 1974)). Appellants believe that the cited art fails to meet this standard. For the reasons above, independent claim 1, 13 and 21 are allowable over the cited art. In addition, dependent claims 2-4, 7, 8, 10, 12, 14-18, 22-25 and 27 are allowable because each depends upon one of the allowable independent claims. Therefore, Appellants submit that the current grounds of rejection are in error and respectfully request the §103(a) rejections of claims 1-3, 7, 8, 10, 12-18, 21-25 and 27 be reversed.

Respectfully submitted,

Date: July 16, 2009

/Gregory K. Goshorn/

By: Gregory K. Goshorn

Reg. No.: 44,721

ATTORNEY FOR APPELLANTS

Greg Goshorn, P.C.
9600 Escarpment
Suite 745-9
Austin, Texas 78749
Telephone: (512) 291-9203
Facsimile: (512) 535-4206

CLAIMS APPENDIX

(Currently Pending Claims)

1. (Previously presented) A method for predicting service level in a utility computing environment having a dynamically allocated subset of computing resources from a set of available computing resources, the method comprising:

creating a resource profile corresponding to the dynamically allocated subset of computing resources allocated according to a service level agreement;

loading a workload profile representing a demand profile for an enterprise;

simulating processing of the workload profile using the resource profile to produce a service level result, wherein the resource profile subset is modified during the simulation according to the service level agreement;

and

generating a new service level agreement in the event the resource profile cannot process the workload profile at an expected service level corresponding to the service level agreement, wherein the new service level agreement will process the workload profile at the expected service level.

2. (Previously presented) The method of claim 1, further comprising:

comparing the service level result to the service level agreement; and

signaling whether the resource profile will process the workload profile at an expected service level corresponding to the service level agreement.

3. (Previously presented) The method of claim 1, wherein the subset of computing resources includes allocated processing resources and memory resources for a client account.

4. (Previously presented) The method of claim 1, wherein the service level agreement includes a base resource allocation, a maximum resource allocation, resource costs, and rules for dynamically reallocating the resources based upon workload demand.

5. (Previously presented) The method of claim 1, wherein the simulation is scheduled to run automatically at an off-peak time.

6. (Previously presented) The method of claim 1, further comprising determining a cost associated with meeting the new service level agreement.

7. (Previously presented) The method of claim 1, wherein the resource profile includes a communication bandwidth allocation.

8. (Previously presented) The method of claim 1, further comprising comparing the workload profile to a second workload profile representing an actual demand profile for a second client account;

wherein the simulating is based upon a result of the comparison.

9. (Cancelled)

10. (Original) The method of claim 1, wherein the workload profile includes scheduling information and the simulation step incorporates the scheduling information in the processing.

11. (Original) The method of claim 1, wherein the workload profile includes information corresponding to one or both of prioritization of resources and importance of specific resources.

12. (Original) The method of claim 1, wherein the workload profile is loaded from a configuration file.

13. (Previously presented) A system for simulating service in a utility computing environment having a first service level agreement to service the demands of an enterprise using a dynamically allocated subset of computing resources from a set of available computing resources, comprising:

an allocated subset of the set of computing resources;

logic for loading a workload profile representing a hypothetical demand profile for a client account;

logic for simulating processing of the workload profile, wherein the workload profile is based upon actual, measured data, using the allocated subset of the set of available computing resources to produce a service level result;

logic for modifying the allocated subset of the available computing resources

based upon the service level result; and

logic for generating a new service level agreement in the event the simulation

produced by the simulation logic cannot process the workload profile at an

expected service level corresponding to the first service level agreement,

wherein the new service level agreement will process the workload profile

at the expected service level.

14. (Previously presented) The system of claim 13, further comprising:

logic for comparing the service level result to a second service level agreement;

and

logic for signaling whether or not modified, allocated subset of the available computing resources will process the workload profile at an expected service level corresponding to the second service level agreement.

15. (Previously presented) The system of claim 13, wherein the set of available computing resources comprises:

processing resources; and

memory resources.

16. (Previously presented) The system of claim 15, wherein the set of available computing resources further comprises:

a base resource allocation;

a maximum resource allocation;
resource costs; and
rules for dynamically reallocating the resources based upon workload demand.

17. (Previously presented) The system of claim 15, wherein the set of available computing resources further comprises communication bandwidth.

18. (Original) The system of claim 13, further comprising logic for comparing the workload profile to a second workload profile representing an actual demand profile for a second client account;

wherein a simulation produced by the simulation logic is based upon a result of the comparison step.

19. (Canceled)

20. (Original) The system of claim 13, wherein the workload profile includes scheduling information and the simulation logic incorporates the scheduling information in the processing.

21. (Previously presented) A computer program product for predicting service level compliance in a utility computing environment having a service level agreement to service the demands of an enterprise using a dynamically allocated subset of computing resources from a set of available computing resources, comprising:

a memory;

a resource list, stored on the memory for execution on a processor, detailing a set of available computing resources;

an allocated resource list, stored on the memory, detailing an allocated subset of the set of available computing resources;

logic, stored on the memory for execution on a processor, for creating a computer resource profile based upon the allocated subset of the set of available computing resources;

logic, stored on the memory for execution on a processor, for loading a workload profile representing a hypothetical demand profile for a client account;

logic, stored on the memory for execution on a processor, for simulating the processing of the workload profile using the computer resource profile to produce a service level result;

logic, stored on the memory for execution on a processor, for comparing the service level result to a service level agreement;

logic, stored on the memory for execution on a processor, for signaling whether the computing resource profile will process the workload profile at an expected service level corresponding to the service level agreement; and

logic, stored on the memory for execution on a processor, for generating a new service agreement in the event the computing resource profile cannot process the workload profile at the expected service level corresponding to the service level agreement, wherein the new service level agreement will process the workload profile at the expected service level.

22. (Previously presented) The system of claim 21, wherein the computing resource profile comprises:

processing resources; and

memory resources.

23. (Original) The system of claim 22, wherein the computing resource profile further comprises:

a base resource allocation;

a maximum resource allocation;

resource costs; and

rules for dynamically reallocating the resources based upon workload demand.

24. (Previously presented) The system of claim 22, wherein the computing resource profile also comprises a communication bandwidth allocation.

25. (Original) The system of claim 21, further comprising logic for comparing the workload profile to a second workload profile representing an actual demand profile for a second client account;

wherein a simulation produced by the simulation logic is based upon a result of the comparison step.

26. (Canceled)

27. (Original) The system of claim 21, wherein the workload profile includes scheduling information and the simulation logic incorporates the scheduling information in the processing.

EVIDENCE APPENDIX

No evidence has been submitted in conjunction with this application.

RELATED PROCEEDINGS APPENDIX

There are currently no related proceedings associated with this application.